

Appl. No. 10/669,822
Amdt. dated June 23, 2006
Reply to Office action of April 7, 2006

REMARKS

Applicant has received the Office action dated April 7, 2006, in which the Examiner: 1) objected to the Abstract; 2) rejected claim 1 as allegedly failing to comply with the written description requirement; 2) rejected claim 3 as allegedly indefinite; 3) rejected claim 1 as allegedly directed to non-statutory subject matter; 4) rejected claims 15-18 as allegedly anticipated by Mikesell (U.S. Pub. No. 2004/0153479, hereinafter "Mikesell"); 5) rejected claims 8-14 and 19-25 as allegedly unpatentable over Cannon (U.S. Pat. No. 5,983,239, hereinafter "Cannon"); and 6) rejected claims 1-7 as allegedly unpatentable over Cannon in view of Howard (U.S. Pat. No. 6,519,612, hereinafter "Howard").

With this Response, Applicant amends claims 1-8 and 19. Reconsideration is respectfully requested.

I. OBJECTIONS TO THE ABSTRACT

With this Response, Applicant amends the Abstract to address the Examiner's concern. No new matter is added.

II. 35 U.S.C. § 112 REJECTIONS

The Office action of April 7, 2006 rejects claim 1 as allegedly failing to comply with the written description requirement regarding the "autonomous of the user" limitation. Applicant respectfully traverses.

First, it is noted that "autonomous of the user" limitation is not a global limitation regarding claim 1, but is rather a further refinement of the implementing limitation. Thus, it follows that the user can partake in other limitations, e.g., creating metadata. Implementing a storage strategy "autonomous of the user" is discussed throughout the specification. As an example, the Examiner's attention is drawn to Applicant's Figure 4 and related disclosure.

In a related rejection, the Office action rejects claim 3 as indefinite, citing an allegedly contradiction between claim 3 and claim 1. Again, Applicant respectfully traverses. As noted immediately above, the "autonomous of the user" is a further refinement of the implementing limitation, and does not affect the creating limitation. Thus, there is no contradiction between claim 3 and claim 1.

Appl. No. 10/669,822
Amdt. dated June 23, 2006
Reply to Office action of April 7, 2006

Based on the foregoing, Applicant respectfully submits that claim 1 does not suffer from written description shortcomings, that claim 3 does not suffer from indefiniteness, and the rejections should be withdrawn.

III. 35 U.S.C. § 101 REJECTIONS

The Office action rejects claim 1 as allegedly directed to non-statutory subject matter. Applicant amends claim 1 and its dependent claims to be directed to a computer readable medium storing a program to that implements the method to address the rejection.

IV. ART-BASED REJECTIONS

A. Claim 1

Claim 1 stands rejected as allegedly obvious over Cannon and Howard. Applicant amends claim 8 to more clearly define over the file aggregation system of Cannon, and as discussed with respect to the Section 101 rejection.

Cannon is directed to a storage management system with file aggregation supporting multiple aggregated file counterparts. (Cannon Title). In particular, Cannon discusses the existence of processing overhead associated with a file in a system, and that for smaller files the processing overhead is a predominant factor in accessing the smaller files.

The storage of each file requires both media preparation overhead and bookkeeping overhead, delaying completion of the entire storage process. The overhead for storage of a file is independent of that file's size. Thus, the overhead for a large file is overshadowed by its more substantial I/O time. The opposite is true with small files, where the necessary overhead dominates the file storage process compared to the file's relatively short I/O time. Consequently, I/O time is the chief obstacle in speedier storage of large files, whereas overhead prevents small files from being stored faster.

(Cannon Col. 1, line 65 through Col. 2, line 7). In order to decrease the effect of the processing overhead associated with each access to a file, Cannon discloses a data storage subsystem 102 that performs services for client stations 106 (e.g., archive and backup) where user files are aggregated into managed files.

One of the key features of the present invention is storage and use of **"managed" files, each comprising an aggregation of one or multiple constituent "user" files.** The "user" files are created by

Appl. No. 10/669,822
Amdt. dated June 23, 2006
Reply to Office action of April 7, 2006

BEST AVAILABLE COPY

the client stations 106, and managed by the subsystem 102 as a service to the client stations 106. This "internal" management scheme [of the subsystem 102] helps to significantly reduce file management overhead costs by using managed files constructed as aggregations of many different user files. In particular, **the subsystem 102 treats each managed file as a single file during backup, move, and other subsystem operations, reducing the file management overhead to that of a single file.**

(Cannon Col. 7, lines 54-67 (emphasis added)). Cannon discloses tables that track the membership of various user files in a managed file; however, the user files appear are consistently identified by the name assigned by client stations 106. (Cannon Col. 8, lines 29 through Col. 9, line 11; Table 1 (note "user file name" column); Table 3 (note "user file name" column)).

Claim 1, by contrast, specifically recites, "receiving a file from a client machine, the receiving by appearing to operate in the client machine namespace and in the client machine file structure; ... implementing, autonomously of a user of the file, storage strategies for the file based on the metadata and in a namespace different than the client machine name space." Applicant respectfully submits that Cannon and Howard do not teach or suggest such a system. In particular, Cannon appears to teach tracking a file throughout the Cannon system by the file name assigned by the client station. Howard is cited only for "implementing storage strategies autonomously of the user." (Office action of April 7, 2006, page 9, numbered paragraph 9). Thus, even if hypothetically the teachings of Howard are precisely as the Office actions suggests, (which Applicant does not admit is proper), Cannon and Howard still fail to teach or suggest "implementing ... in a namespace different than the client machine name space."

Based on the forgoing, Applicant respectfully submits that claim 1, and all claims which depend from claim 1 (claims 2-7), should be allowed. Applicant amends claims 2-7 to reflect the amendment regarding a computer readable medium, and not to define over any cited art.

Appl. No. 10/669,822
Amdt. dated June 23, 2006
Reply to Office action of April 7, 2006

BEST AVAILABLE COPY

B. Claim 8

Claim 8 stands rejected as allegedly obvious over Cannon. Applicant amends claim 8 to more clearly define over the file aggregation system of Cannon.

Cannon is directed to a storage management system with file aggregation supporting multiple aggregated file counterparts. (Cannon Title). In particular, Cannon discusses the existence of processing overhead associated with a file in a system, and that for smaller files the processing overhead is a predominant factor in accessing the smaller files. (Cannon Col. 1, line 65 through Col. 2, line 7). In order to decrease the effect of the processing overhead associated with each access to a file, Cannon discloses a data storage subsystem 102 that performs services for client stations 106 (e.g., archive and backup) where user files are aggregated into managed files. (Cannon Col. 7, lines 54-67). Cannon discloses tables that track the membership of various user files in a managed file; however, the user files appear to consistently be identified by the name assigned by client stations 106. (Cannon Col. 8, lines 29 through Col. 9, line 11; Table 1 (note "user file name" column); Table 3 (note "user file name" column)).

Claim 8, by contrast, specifically recites, "wherein the host computer communicates files to the server for storage on at least one of the plurality of storage devices, wherein the server appears to be a network storage device operating in a user name space and in a user file structure; and wherein a program executing on the server selects on which of the plurality of storage devices to store the files on a file-by-file basis based on storage characteristic preferences supplied for each file, and wherein each file is stored under a globally unique name in a global namespace of the server." Applicant respectfully submits that Cannon does not teach or suggest such a system. In particular, Cannon appears to teach tracking a file throughout the Cannon system by the file name assigned by the client station. Thus, Cannon does not teach or suggest "and wherein each file is stored under a globally unique name in a global namespace of the server."

Appl. No. 10/669,822
Amdt. dated June 23, 2006
Reply to Office action of April 7, 2006

BEST AVAILABLE COPY

Based on the forgoing, Applicant respectfully submits that claim 8, and all claims which depend from claim 8 (claims 9-14), should be allowed.

C. Claim 15

Claim 15 stands rejected as allegedly anticipated by Mikesell.

Mikesell is directed to systems and methods for restriping files in a distributed file system. (Mikesell Title). In particular, Mikesell is directed to a distributed file system 110 comprising multiple "smart" storage units 114. (Mikesell Paragraph [0062]; Figure 1). Data files are stored and accessed as a standard file system, while blocks of each file may be spread across the multiple storage units.

The systems and methods of the present invention provide an intelligent distributed file system, **which enables the storing of data among a set of smart storage units that are accessed as a single file system.**

(Mikesell Paragraph [0045] (emphasis added)).

The intelligent distributed file system 110 enables blocks of an individual file to be spread across multiple smart storage units 114.

(Mikesell Paragraph [0062]). The Mikesell system keeps metadata about each file that identifies the storage location of each block of data that makes up the file.

The intelligent distributed file system advantageously utilizes a metadata data structure to track and manage detailed information about each file, including, for example, the device and block locations of the file's data blocks.

(Mikesell Paragraph [0008]).

Claim 15, by contrast, specifically recites, "wherein the server appears to programs executing on the client computer as a network storage device operating in a user namespace and in a user file structure; and wherein the server stores the file on at least one of the first and second storage devices in a global namespace" Applicant respectfully submits that Mikesell does not expressly or inherently teach such a system. In Mikesell, data files are stored and accessed as a standard file system. Thus, Mikesell does not expressly or inherently teach "wherein the server appears to programs executing on the client

Appl. No. 10/669,822
Amdt. dated June 23, 2006
Reply to Office action of April 7, 2006

BEST AVAILABLE COPY

computer as a network storage device **operating in a user namespace and in a user file structure; and wherein the server stores the file on at least one of the first and second storage devices in a global namespace ...**."

Based on the foregoing, Applicant respectfully submits that claim 15, and all claims which depend from claim 15 (claims 16-18), should be allowed.

D. Claim 19

Claim 19 stands rejected as allegedly obvious over Cannon. Applicant amends claim 19 to more clearly define over the file aggregation system of Cannon.

Cannon is directed to a storage management system with file aggregation supporting multiple aggregated file counterparts. (Cannon Title). In particular, Cannon discusses the existence of processing overhead associated with a file in a system, and that for smaller files the processing overhead is a predominant factor in accessing the smaller files. (Cannon Col. 1, line 65 through Col. 2, line 7). In order to decrease the effect of the processing overhead associated with each access to a file, Cannon discloses a data storage subsystem 102 that performs services for client stations 106 (e.g., archive and backup) where user files are aggregated into managed files. (Cannon Col. 7, lines 54-67). Cannon discloses tables that track the membership of various user files in a managed file; however, the user files are consistently identified by the name assigned by client stations 106. (Cannon Col. 8, lines 29 through Col. 9, line 11; Table 1 (note "user file name" column); Table 3 (note "user file name" column)).

Claim 19, by contrast, specifically recites, "wherein the first means for executing communicates files to the second means for executing for storage on at least one of the plurality means for storing, wherein the second means for storing to be a network storage device operating in a file structure of the first means for executing; and wherein program executing on the second means executing selects on which of the plurality of means for storing to store the files on a file-by-file basis based on storage characteristic preferences supplied for each file, and wherein each file is stored under a globally unique name in a global namespace of the plurality of means for storing." Applicant respectfully submits that Cannon

Appl. No. 10/669,822
Amdt. dated June 23, 2006
Reply to Office action of April 7, 2006

BEST AVAILABLE COPY

does not teach or suggest such a system. In particular, Cannon appears to teach tracking a file throughout the Cannon system by the file name assigned by the client station. Thus, Cannon does not teach or suggest "and wherein each file is stored under a globally unique name in a global namespace of the plurality of means for storing."

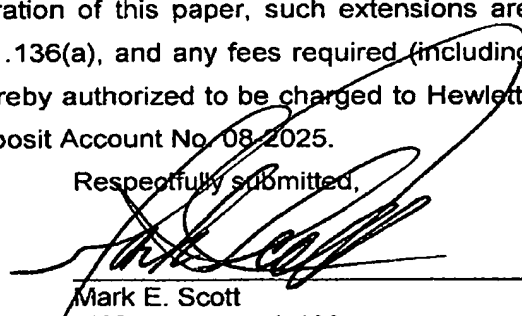
Based on the forgoing, Applicant respectfully submits that claim 19, and all claims which depend from claim 19 (claims 20-25), should be allowed.

V. CONCLUSION

In the course of the foregoing discussions, Applicant may have at times referred to claim limitations in shorthand fashion, or may have focused on a particular claim element. This discussion should not be interpreted to mean that the other limitations can be ignored or dismissed. The claims must be viewed as a whole, and each limitation of the claims must be considered when determining the patentability of the claims. Moreover, it should be understood that there may be other distinctions between the claims and the cited art which have yet to be raised, but which may be raised in the future.

Applicant respectfully requests reconsideration and that a timely Notice of Allowance be issued in this case. It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's Deposit Account No. 08-2025.

Respectfully submitted,



HEWLETT-PACKARD COMPANY
Intellectual Property Administration
Legal Dept., M/S 35
P.O. Box 272400
Fort Collins, CO 80527-2400

Mark E. Scott
PTO Reg. No. 43,100
CONLEY ROSE, P.C.
(512) 391-1900 (Phone)
(512) 320-9181 (Fax)
ATTORNEY FOR APPLICANT